Topography from Shading and Stereo

Berthold K.P. Horn
Artificial Intelligence Laboratory, Massachusetts Institute of Technology
Michael Caplinger
Department of Geology, Arizona State University

Abstract: *Methods exploiting photometric information in images that have been developed in machine vision can be applied to planetary imagery. Present techniques, however, focus on one visual cue, such as shading or binocular stereo, and produce results that are either not very accurate in an absolute sense or provide information only at a few points on the surface. We plan to integrate shape from shading, binocular stereo and photometric stereo to yield a robust system for recovering detailed surface shape and surface reflectance information. Such a system will be useful in producing quantitative information from the vast volume of imagery being received, as well as in helping visualize the underlying surface. The work will be carried out on a popular computing platform so that it will be easily accessible to other workers.*

1. Introduction

It would be very useful for users of image data to have automated means of extracting accurate topography from images. Existing automated methods are not able to robustly recover detailed surface shape. We propose to explore the intimate integration of existing machine vision methods to recover topography, and possibly also surface reflectance. This will be a demonstration of the application of a particular machine vision paradigm to problems of data reduction and visualization in the space sciences.

Present binocular stereo methods, whether correlation or edge-based, cannot deal with large differences in foreshortening and lighting between the two images. Conversely, so-called shape-from-shading methods cannot accurately recover the lower spatial frequency components of the topography, and may be misled by spatial variations of surface reflectance[1]. The two methods are complementary in that binocular stereo can provide sparse absolute height data, while shading provides fine detail. Furthermore, binocular stereo cannot provide reliable information in areas where there is little texture, while shading works best where the surface curves smoothly and has near uniform reflectance properties.

[1] What we call reflectance variations here are often referred to as albedo variations—we avoid the term "albedo," since some people at least consider this to be a technical term with a meaning different from the one intended here.

Recent progress in shape from shading, multi-resolution grey-level matching binocular stereo, as well as photometric stereo, suggests that robust methods may be designed that combine the available information to recover accurate surface shape and reflectance information. This "fusion" of early vision modules cannot, however, be done at a point where each module has already generated its error-prone output. Instead, the available cues must be used in an integrated way. The now classic calculus of variations approach to solving early vision problems provides a means for achieving this synergism.

We will explore this new approach to integration of early vision modules in the context of the interpretation of multiple images of the same surface area obtained from different viewpoints, possibly under different lighting conditions. The result will be a system that can recover accurate, detailed topographic information, as well as surface reflectance. This will greatly enhance the value of the voluminous image data now being returned to earth from cameras on planetary explorer spacecraft, as well as earth observation platforms and from cometary flybys.

2. Applications

The volume of imaging data being received is growing all the time, and there is no hope of extracting topography from it via manual or semi-automated methods except on a piecemeal basis.

Geophysical uses of topographic data, both on Earth and other planets, typically involve the study of the movement of material over the surface under the influence of gravity. Such movement is both influenced by the existing topography and the cause of topographic changes (as in erosion). The flowing material can be water or other liquids, rock in mass movement, lava, ice, the atmosphere, or mixtures of these.

- **Hydrology:** For example, for the flow of water, the velocity of a flow that created a channel can be calculated using empirical relationships relating variables such as slope and cross-sectional area. The longitudinal profile of a stream (the slope of the fluid bed as a function of position downstream) reflects the nature of fluid processes occurring at the bed and is a sensitive measure of both surface and sub-surface sources of fluid. Such relationships have been used to explore both terrestrial water erosion and the possible genesis by water of the martian channels.
- **Mass movements:** Mass movements, such as landslides and avalanches, are caused when the gravitational force on a body of material overcomes the forces (such as shear strength) that bind it. (Sometimes, gravity is assisted by accelerations induced by earthquakes or other vibrations.)

After such movement, the volume of material moved can be computed directly from the topography, and its speed and rheological properties estimated from slope relationships and the relief of specific features on the surface of the flow.

- **Volcanism:** Lava also flows under the influence of gravity. Topographic data can be used to determine the volume of extrusion and intrusion (important in constraining processes that occur in the magma chamber), as well as to determine rheological properties of flows from the appearance of surface features. These rheological properties are in turn related to chemical and mineralogical composition, eruption rate, temperature, and other physical parameters of the lava.
- **Glaciology:** Ice, too, moves under its own weight. The slope of a ice-stream's surface reflects both the underlying topography and the rheological behavior of the ice, in particular its velocity and volume discharge rate. The relief of surface features also permits analysis of the ice depth where it cannot be directly measured.
- **Aeolian processes:** The flow of the atmosphere over topographic features is both controlled by their shape and contributes to their erosion and modification, especially if solid material is carried by the atmosphere in saltation or suspension. Knowledge of the topography can be used to estimate the nature of the flow (that is, turbulent or laminar) and the nature of material being carried by wind.

Other applications of topographic data, not related to geology directly, can also be mentioned. Some are: (a) mission planning, both on the surface and from aircraft and low-altitude satellites; (b) meteorological modeling: to constrain the surface wind field and model the generation of turbulence. Note, by the way, that much of the above applies to observations of earth also, not only to other planets.

3. Background

A considerable fraction of the bits returned to earth by spacecraft sent to explore the planets comes in the form of images. Presently the means to obtain quantitative information from these images are largely restricted to photogrammetric methods and profile-based photoclinometric methods. Yet a human observer gets a wealth of additional qualitative information from these images that suggest that it ought to be possible to extract more, using advanced machine vision techniques. If detailed surface shape and reflectance information can be recovered, it can be presented in a variety of ways to aid in visualization using well-known rendering methods from computer graphics.

We want to exploit this opportunity by applying the latest methods in work on "early vision." First we review traditional methods of image analysis.

3.1 Photogrammetry

Photogrammetric methods permit the accurate determination of the positions of isolated points in three-dimensional space from (typically manual) spot measurements of corresponding images of well defined features on the surface. This is a very important basic step in defining the surface, and in recovering the camera position and attitude, but provides information at only a small number of isolated points, which do not define the detailed shape in between. Nevertheless, sound photogrammetric techniques are vital to the determination of an accurate reference body for use in cartography and in determining the relative positions and orientations of exposure stations used to obtain image pairs for binocular stereo [Horn 89].

3.2 Photoclinometry

Present photoclinometric methods permit the recovery of isolated profiles across features that have special symmetries, such as circular craters, volcanic calderas, and linear depressions or grabens [Bonner & Schmall 73] [Davis *et al.* 82] [Davis & Soderblom 83] [Davis & McEwen 84] [Howard *et al.* 82] [Lambiotte & Taylor 67] [Lucchitta & Gambell 70] [Malin & Danielson 84] [McEwen 85] [Passey & Shoemaker 82] [Rowan *et al.* 71] [Tyler *et al.* 71] [Watson 68] [Wildey 86] [Wilhelms 64] [Wilson *et al.* 84] (for a larger collection of references on this subject, see [Horn & Brooks 90]). Since these methods cannot take into account cross-profile inclination, they will not produce accurate profiles if the cross-profile inclination is non-zero. Hence such methods are limited to areas that have the appropriate symmetry. Shape-from-shading, basic to what is being proposed here, may be thought of as "area-based" photoclinometry. It permits the recovery of complex, wrinkled surface shapes by using image information from a full two-dimensional region of the image, rather than merely along a line in the image. There is no restriction to surfaces having predefined symmetries (see section 4.2).

3.3 Correlation-based Stereo

There have been many attempts to automate the recovery of topographic information from two images using binocular stereo methods. Perhaps the oldest and most widely known are methods based on correlation. A wide-variety of techniques, ranging from Fourier transforms to optical computing have been employed to speed the computation [Horn 86]. Such methods assume that, locally at least, what appears in one image is a shifted version of what appears in the other image. If, however, the surface is viewed from quite different viewpoints, the foreshortening in the two images will be very different and these methods will fail, since they are not designed to match two waveforms that have been stretched by different amounts [Horn 83]. Such methods are thus restricted to situations where the baseline-to-height ratio is small. Unfortunately, the expected error in the determination of the distance from the camera to the surface grows inversely as the baseline-to-height ratio—so these methods have not proven useful in the interpretation of aerial photographs, for example, where a large baseline-to-height ratio is purposefully employed to get accuracy in the determination of height comparable to the accuracy in the determination of the horizontal position. Furthermore, in planetary exploration it is common to have two images of the same area taken not only from very different viewpoints, but under very different lighting conditions. Correlation-based binocular stereo is based on the assumption that the patterns of brightness in the two images are the same and so cannot be applied successfully in this situation.

Of the dozens of attempts over the past twenty years to automate binocular stereo using this kind of approach, apparently the only one that comes close to being useful is embodied in the so-called Gestalt photomapper. On undulating terrain with sufficient texture and no confusing reflections from specular surfaces such as lakes, this machine can produce beautifully detailed surface topography. It does, however, require considerable assistance from an experienced operator in order to help it out in difficult areas. It also does not work satisfactorily in areas of steep relief or where there is insufficient contrast in texture patterns. While the detailed inner workings of the machine are proprietary, it is known that part of the reason that it works at all is that it is not based on blind correlation, but an iterative scheme that warps the local surface in a hexagonal patch using a high-order polynomial. In this fashion it can take account of the differences in foreshortening, provided that the initial guess is close enough to the correct shape to allow it to converge.

3.4 Edge-Based Stereo

To overcome the problems of unequal foreshortening and unequal brightness, there has been considerable work on feature-based binocular stereo in the last ten years. In this case distinctive "features" of the images (such as rapid transitions in grey level) are first extracted and then these features are matched symbolically. The matching process is complex and requires representation of the images at multiple resolutions. It does not lend itself readily to parallel implementation.

The basic assumption in feature-based binocular stereo is that image features correspond to distinguished surface features (such as terrain breaks) and that their image position is not affected by vagaries of lighting and foreshortening. With low baseline-to-height ratio and similar lighting in the two images, such methods can produce reasonable estimates of height along isolated curves. The remaining surface is unknown and has to be somehow interpolated from the sparse data recovered [Grimson 82]. This means that in areas where there are few "features" very little is really known about local surface topography.

4. Photometric Machine Vision Methods

In machine vision, a number of methods have been developed that exploit the fact that a grey-level is a quantized estimate of image irradiance, and that image irradiance depends on surface orientation, surface reflectance and light source distribution. Such methods have shown promise in applications to aerial photographs, satellite photographs of the earth and planetary images. We briefly review the relevant approaches here, noting that what we are proposing now is an integration of these methods—but not one that merely builds a system out of several existing modules. We are not proposing to merely combine the outputs of various modules in some simple statistical way. Instead, the calculus of variations approach is to be used to minimize some overall error function. We believe that this is the way to obtain the necessary synergism between different visual cues.

4.1 Photometric Stereo

Images of an object taken from the same viewpoint but under different lighting conditions can be used to recover surface shape. This is called *photometric stereo*, in distinction to *binocular stereo* discussed above [Horn *et al.* 78] [Horn 86]. The basic idea is that local surface orientation on a more or less smooth surface can be specified using two parameters (such as the slope in the x-direction and the slope in the y-direction). A measurement of brightness at the corresponding image point places a single constraint on the surface orientation—not enough to recover it fully. To recover surface orientation, various other sources of additional constraint can be exploited. Perhaps the easiest is the brightness, at the same point in an image, taken when the object is illuminated differently.

With two images, two constraints are available—that is, it becomes possible to locally recover the surface orientation. The computation can even be implemented in a lookup table indexed on the two grey-levels at the same picture cell. The lookup table can be constructed using some theoretical photometric model and known illumination conditions, or it can be developed by numerical inverting calibration data obtained from an object of known shape.

One can do even better if more than two images are available. If, for example, the surface reflectance is spatially varying, then two images are not sufficient to locally recover the three degrees of freedom, but three images will do. Continuing in this way, if the photometric model has n unknown parameters that may vary from point to point, then $(n + 2)$ images are needed to estimate these parameters, as well as the two components of surface orientation[2].

4.2 Shape from Shading

Shape from shading is the recovery of surface shape from shading in a single image. Shading is the spatial variation of image brightness resulting from corresponding variations in surface orientations. The twenty year history of work in this area is captured in a recently published collection of papers called *Shape from Shading* [Horn & Brooks 89]. This book also contains a complete bibliography of the three hundred or so publications in this, and related fields, such as photoclinometry—so we will keep the list of references short here. For a quick introduction to the subject, see chapters 10 & 11 in *Robot Vision* [Horn 86].

In shape from shading the local ambiguity in surface orientation, described above, is resolved by assuming that the surface "hangs together," that is, neighboring surface patches are not allowed to have totally unrelated orientations. Put another way, if one were to walk in a closed path atop the surface patches one ought to come back to the same height. Use of this additional constraint is, however, much more difficult than use of information from a second image as in photometric stereo. It has taken a long time for theory and implementation to mature to where these methods are really practical on other than synthetic images and real images of relatively simple shapes[3].

[2] There are limitations to this in the context of planetary imaging, since the light source, at least over a short period of time, moves along a small circle on the sphere of directions, and thus does not sample the space of all possible source directions well.

[3] Even at this stage, question of existence and uniqueness of the solution are re-

Nevertheless, robust methods that work on real images of complex, wrinkled surface have been developed in recent years [Frankot & Chellapa 88] [Szeliski 90] [Zheng & Chellapa 90]. An example of what can be done is the surface recovered from a SPOT image of a hilly area east of Huntsville, Alabama by the new method in [Horn 90] (See scetion 8 and figures attached at the end). A reasonable surface shape was recovered despite the fact that the range of distinct grey-levels in the region of interest was only 19, that the photometric function, the light-source position and the atmospheric state were not accurately known, and that the image was noisy and appeared to suffer from the effects of aliasing. The recovered shape did not reproduce the lower spatial frequencies of the terrain very accurately, as expected. Nevertheless, a contour map obtained from a smoothed version of the surface looks similar to the appropriate part of the USGS topographic map of the same area. Conversely, fine surface undulations were recovered that do now show up in the "generalized" topographic map.

Perhaps the earliest work in this field can be traced back to [Rindfleisch 66], who found that very specialized assumptions about the surface reflectance properties (namely that brightness is a linear combination of the slope in the x-direction and the slope in the y-direction) allows one to estimate a profile of the surface by simple integration. Some recent work also uses this assumption, which dramatically simplifies the problem [Pentland 88]. But this is not a realistic assumption, nor is it desirable [Horn 70]. The reason is that the recovered profiles are not connected to one another, each can independently float up or down with respect to its neighbor, so that there is tremendous ambiguity in the recovered surface.

If surface slopes are small, then reflectance properties may be locally linearized and the above used as an approximation. It has been found that more general purpose shape-from-shading methods work particularly well when the slope excursion are small [Kirk 84, 87].

Shape-from-shading suffers from some of the same limitation as photoclinometry, since it is also based on an assumed known photometric function. However, since information from a two-dimensional region is used in a least-squares approach, the effects of image noise and errors in photometry tend to be suppressed. Conversely, photoclinometry is limited because varying surface reflectance along a profile cannot be distinguished from changes in brightness due to surface orientation. When a two-dimensional patch is considered, on the other hand, such variations lead to inconsistencies that can be detected and at least used to flag the area as suspect. Hopefully we will learn more

ceiving considerable attention—as are issues of what kind of boundary conditions might be needed. These do not concern us very much here, since we know that there is a solution, and we can often get a good first estimate of it.

about this in future to actually be able to recover both shape and surface reflectance—although at this stage this does not look too promising since the problem appears to be underconstrained when working from a single image.

4.3 Grey-level Matching Stereo

People have little difficulty interpreting stereo pairs of smooth surfaces lacking the obvious "features" needed by present feature-based binocular stereo methods. As a result, there has been considerable interest in new methods that have been developed to directly match grey-levels, or information derived directly from grey-levels using local operators [Gennert 86] [Barnard 89]. Such methods hold the promise of providing dense surface information, although they have to overcome the problems inherent in the assumption that a given point in the scene will yield similar grey-levels in the two images. It is known that such methods cannot work on a single level of resolution or they will immediately get locked into a local minimum resulting from inappropriate matches between the two images. Multi-resolution techniques are needed to solve this problem.

5. Proposed New Method

There has been quite a bit of work recently on sensor "fusion" and early vision module "integration," driven by the lack of robustness of individual modules. However, most of this work either uses the results of one module to constraint the operation of another, or simply combines the results of several module using some simple least-squares weighting function. Neither of these approaches has proven particularly useful. We propose instead to integrate the modules much more intimately. Over the past ten years, a new approach to machine vision has arisen, based on the calculus of variations, as pioneered in [Horn & Schunck 81] and [Ikeuchi & Horn 81]. Typically an error functional is defined that has several component penalty functions that allow one to express a preference for solutions having certain properties, such as smoothness. This helps one deal with what would otherwise be underconstrained or ill-conditioned problems.

What we propose to do is to build an energy functional that contains penalty terms for errors in shading in the left image, errors in shading in the right image and errors in matching left and right image. We will then have to find the equations governing the resulting variational problem and develop methods for iteratively obtaining a solution to a discrete version of the resulting partial differential equations. The hope is that the problem can be formulated in such a way that the not inconsiderable investment in work on the latest shape-from-shading system [Horn 90] can be somehow carried over into the new system for shape from shading and binocular stereo.

Initially the main concern will not be computational efficiency, except in so far as to avoid methods that will not later lend themselves to speedup using numerical approaches such as multigrid and gradient descent methods, or massively parallel hardware such as the Connection Machine. On a low-end work-station we do not expect to be able to process large images in a short time. The focus will be on developing a sound mathematical approach first. We do plan, however, to incorporate, as soon as possible, ideas on high speed implementations of the iterative algorithms being pursued by [Szeliski 90] [Zheng & Chellapa 90].

The resulting system will be tested both on synthetic image pairs and real image pairs. The reason for working with synthetic image data first is that this is the only situation in which "ground truth" is known absolutely. But of course a system that only works on clean, noise-free synthetic data is of little interest, so test will have to be made on real data also. We propose to do this first with satellite images of the surface of the earth, again because the "ground truth" is more readily accessible to us in this case. Finally, the system will be tested on existing planetary images.

5.1 Limitations on Accuracy

As mentioned before, shape-from-shading methods can provide fine detail, but not accurately reproduce the lower spatial frequency undulations of the terrain. Conversely, binocular stereo can give good spot absolute height measurements, but cannot provide dense coverage of the surface, particularly in areas without sufficient texture. We expect that the integration of these two modalities will overcome the short-comings of both method. There will, however, still be limitations resulting from short-comings of the imaging systems and limited knowledge of surface photometry.

In profile-based photoclinometry, the effects of errors in photometry are severe, since the measured brightness is directly translated into surface slope. There is no way locally to tell the difference between a change of brightness resulting from a spatial variation in surface reflectance and one resulting from a change in surface orientation. In area-based shape-from-shading, information from a two-dimensional image region is used, and certain errors tend to cancel out. At the very least, the inability to reduce the error functional flags the solution as suspect. Furthermore, when this method is integrated with binocular stereo, it can be expected that the lower spatial frequencies in the reconstruction will be largely controlled by the stereo data and so the effects of errors in photometry will be greatly reduced.

Most of the imagery we propose to work with was acquired using vidicon-based cameras and so may suffer from poor radiometric calibration, non-linear response, spatially non-uniform response and geometric warping. Even though attempts are made to remove these effects through careful pre-processing, we still expect that the results obtained from these images will be inferior to those possible in future with CCD-based cameras.

6. Proposed Work

The proposed work can be divided into six components:

6.1 Application of Photometric Stereo

To become more familiar with photometric properties of planetary surfaces, and to explore the potential for simultaneous analysis of multiple images of the same surface, we propose to start with photometric stereo analysis on images taken with different lighting conditions but from essentially the same exposure station. It will be best at first to work with images obtained from a totally static platform. For this reason we would like to start with certain Viking lander images. We will first assume a photometric function (such as some popular combination of Hapke, Minneart, Lommel-Seeliger, and Lambertian models) and try and recover surface orientation and surface reflectance from three images. We will then work backwards from a calibration object with approximately known surface shape to a photometric function, possibly using many more than three images. Once the lookup table has been constructed it can be used to interpret the rest of the scene. Numerical photometric information obtained, using the same sensor as that used later in measurements of shape, is likely to be much more useful than that provided by an arbitrary analytical form—except in so far as that existing analytical forms may be helpful in bridging gaps in measurements and smoothing out noise.

We realize that the photometric properties of surfaces at different scales are typically very different, so that the detailed results on close range imagery will not apply directly to images taken from orbital distances. Yet the principle involved is the same, and having images taken from *exactly* the same point enables photometric stereo analysis.

6.2 Photometric Stereo and Binocular Stereo

We plan next to investigate the integration of photometric stereo (same position—different lighting) with binocular stereo (same lighting—different position). This promises to be easier to achieve than our ultimate goal of integrating shading and binocular stereo. At the same time we expect to learn valuable lessons from this exercise since some of the same mathematical tools and programming techniques come into play. We expect that even here we will need to develop suitable multi-scale algorithms to overcome the problems of local minima resulting from stereo mismatches.

It should be pointed out that integration of photometric stereo and binocular stereo on close-range image sets may have applications to autonomous vehicle control, since shapes of the surfaces in the environment can be recovered, as well as their surface reflectance patterns. It is quite likely that surface recovery using such integrated methods will involve less complex computations than those from binocular stereo alone[4].

6.3 Application of Shape from Shading

We then propose to apply the latest shape-from-shading method to recovery of shape from existing imagery such as certain Viking Orbiter images. We plan also to work on the more complex problem of "whole disk" shape from shading using images of Deimos. This is more difficult because the boundary constraints on slope are harder to integrate—-since slope becomes infinite on the occluding boundary.

6.4 Shading and Stereo—Similar Lighting, Similar Viewpoint

Next, we come to the heart of the work proposed here, the integration of shape from shading and grey-level matching stereo. We propose to first work on satellite images of the earth, since we have access to reasonably accurate ground truth in this case. We also will initially select situations where the view points are not too wildly different, surface cover is fairly uniform and the lighting is similar in the two images (for example, using a stereo pair of SPOT

[4] On the other hand, for a rapidly moving platform, integration of direct motion vision and binocular stereo is ultimately likely to be more useful still. We do not propose here to work on this, although we expect that a similar approach can be taken to integration of direct motion vision methods and grey-level matching stereo.

images of a hilly tree-covered region east of Huntsville, Alabama)[5]. The main part of the proposed work will be the development of the mathematical model and the computer implementation of the algorithm to solve this problem.

6.5 Shading and Stereo—Different Lighting, Different Viewpoint

Finally, we propose to extend the integrated shading and binocular stereo approach to harder cases, where surface reflectance may vary, viewpoints may be widely spaced and illumination very different in the images. This in essence will require an integration of photometric stereo methods with shape from shading and binocular stereo. We hope that the approach can be extended to certain Voyager images where the photometric situation is more complex than on the rocky planets of the inner solar system.

6.6 More than two images

Once we have formalized the techniques for integrating shape from shading and binocular stereo in two images, we will consider the solution of problems were more than two images taken from different viewpoints under different lighting conditions are available. We expect that the additional constraint provided will make solutions even more robust, provided one can develop a procedure to gets close enough to the solution so convergence is guaranteed. In this situation integration of methods from all three areas: shape from shading, photometric stereo and binocular stereo will be appropriate.

6.7 High Speed Implementation

We expect initially to work with relatively small images or parts of images in order to allow debugging with a reasonable turn around time on a small workstation. Until we have solved the underlying mathematical and numerical problems, and demonstrated reasonable algorithms, we will not place great emphasis on implementations that buy speed in return for complexity. Finally, however, for the results of this work to be useful on high resolution images, attention will have to be directed to methods for reducing the number of iterations required, such as multigrid and gradient descent approaches [Szeliski 90] [Zheng & Chellapa 90].

[5] We may also try to work with NOAA polar orbiter images here, since they provide stereo coverage at considerably lower cost, but expect that the extremely low resolution side-lap stereo coverage will lead to quite unimpressive results.

7. Requirement and Deliverables

To do what we proposed to do above, we will need access to the indicated image sources. To summarize:

- Viking Lander images (for the initial photometric stereo work).
- Viking Orbiter images of the Martian surface, first for application of the existing shape-from-shading algorithm, and later for further testing of the new integrated algorithm.
- Viking Orbiter images of Martian satellites for whole-disk shape from shading.
- SPOT earth images for usable quality integration of shading and stereo.
- Voyager images for work on images with more difficult photometry.

7.1 Hardware and Software

For this work to have a impact on the planetary science community it will have to be available on a widely-used platform—we cannot expect others to re-implement what is going to be a non-trivial software system. Presently the latest shape-from-shading work is implemented in Common Lisp on a Symbolics Lisp Machine. This machine has a software environment very conducive to rapid prototyping and debugging of software, but it is not widely available. We propose to do all future work in the programming language C on a SPARC station from Sun Microsystems. To this end, we propose to acquire the following hardware:

- Graphics SPARC station from Sun Microsystems.
- CDROM drive to read the proposed Voyager image library.
- DAT drive to read Viking Orbiter data and store other image libraries.
- Scanner for images that are not available in machine readable form.

We have file servers, laser printers, standard 9-track tape drives and software in place to support the work we plan to do. We also have a high quality stereo viewer (although no stereo-comparator or other means of accurate image measurement). We do not have a means of high quality hard-copy graphics output, but expect that for demonstration purposes half-tone output on well-tuned laser printers will be adequate.

We will have to acquire at least one SPOT image stereo pair for this work. Travel to one conference per year should be covered so we can present the results of the work. We will also need to cover travel for one trip from MIT to ASU, and one trip from ASU to MIT per year to facilitate detailed collaboration on this project.

7.2 Technology Transfer

By working on a popular platform, using a widely known programming language, we provide for ease of transfer of what we develop to other sites. The photometric stereo program, the new shape-from-shading program and the integrated shading and binocular stereo system will all be made available. If the integrated photometric stereo and binocular stereo system turns out to work well and be useful (other than as a stepping stone), it too will be made available. This proposal does not, however, request resources that would be needed to provide software support beyond documentation supplied with the programs and technical papers describing the algorithms.

8. Illustrations of Previous Related Work

We attach two representative examples of earlier work on the application of shape-from-shading methods to images of interest to planetary and earth scientists.

Shown in Figure 1a are two images of Deimos taken by Viking Orbiter (339B02 & 428B22). In Figure 1b are shown corresponding surface reconstructions using an implementation by Michael Caplinger of an older shape-from-shading algorithm [Horn & Brooks 85]. Brighter areas are closer to the viewer than darker ones.

Shown in Figure 2a is a stereo pair of portions of two satellite images of Monte Sano State park east of Hunstville, Alabama. In Figure 2b is a synthetic stereo pair constructed using the topography recovered by a shape-from-shading algorithm developed recently by Berthold K.P. Horn [Horn 90]. The input to the algorithm was the left sub-image of Figure 2a (the right sub-image was *not* used). Figure 3a is a portion of the USGS topographic map covering the area, while Figure 3b is a contour map created from a smoothed version of the digital terrain model recoverd using shape from shading.

These examples, while based purely on shape from shading, at once illustrate the promise of such approaches and also show some of the short-comings discussed earlier, which we plan to overcome by intimately integrating binocular stereo with shape from shading.

9. Summary

We plan to intimately integrate shape from shading and grey-level matching stereo to obtain robust recovery of surface topography and surface reflectance from multiple images of planetary targets. In preparation for this, we will (a) demonstrate photometric stereo applied to images obtained from the same view point but under different lighting conditions, (b) demonstrate the latest shape-from-shading algorithm on real planetary images, (c) demonstrate the potential for recovering surface topography and surface reflectance from satellite images of the earth and other planets.

The main contribution of this work will be a considerable increase in the value of existing and planned planetary imagery resulting from the ability to get important quantitative information in an automated fashion. Amongst other things, this will enable quantitative analysis and better means of visualizing the data.

10. References

Barnard, S. (1989) "Stochastic Stereo Matching over Scale," *International Journal of Computer Vision*, Vol. 3, pp. 17-32.

Bonner, W.J. & R.A. Schmall (1973) "A Photometric Technique for Determining Planetary Slopes from Orbital Photographs," U.S. Geological Survey Professional Paper 812-A, pp. 1–16.

Davis, P.A., L.A. Soderblom & E.M. Eliason (1982) "Rapid Estimation of Martian Topography from Viking Orbiter Image Photometry," in *Reports on Planetary Geology Program* (1982), NASA Technical Memorandum 85127, pp. 331–332.

Davis, P.A. & L.A. Soderblom (1983) "Rapid Extraction of Relative Topography from Viking Orbiter Images: II. Application to Irregular Topographic Features," in *Reports on Planetary Geology Program* (1983), NASA Technical Memorandum 86246, pp. 29–30.

Davis, P.A. & A.S. McEwen (1984) "Photoclinometry: Analysis of Inherent Errors and Implications for Topographic Measurement," *Lunar and Planetary Science Conference* XV, 12–16 March, pp. 194–195.

Frankot, R.T. & R. Chellapa (1988) "A Method for Enforcing Integrability in Shape from Shading Algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*. Vol. 10, No. 4, pp. 439–451, July. Also *Shape from Shading*, Horn, B.K.P. & M.J. Brooks (eds.), MIT Press, Cambridge, MA..

Gennert, M. A. (1987) "A Computational Framework for Understanding Problems in Stereo Vision," Ph.D. Thesis, Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA.

Grimson, W.E.L. (1982) "A Computational Theory of Visual Surface Interpolation," *Philosophical Transactions of the Royal Society of London B*, Vol. 298, pp. 395–427.

Horn, B.K.P. (1970) "Shape from Shading: a Method for Obtaining the Shape of a Smooth Opaque Object from One View," Technical Report TR-79, Project MAC, MIT, Cambridge, MA. Technical Report TR-232, MIT AI Laboratory, Cambridge, MA. Also in *Shape from Shading*, Horn, B.K.P. & M.J. Brooks (eds.), MIT Press, Cambridge, MA..

Horn, B.K.P. (1983) "Non-Correlation Methods for Stereo Matching," *Photogrammetric Engineering and Remote Sensing*, Vol. 49, No. 4, pp. 535–536, April.

Horn, B.K.P. (1986) *Robot Vision* (particularly Chapters 10 & 11, pp. 202–277), MIT Press, Cambridge, MA & McGraw-Hill, New York, NY.

Horn, B.K.P. (1989) "Relative Orientation," to appear in *International Journal of Computer Vision*, Also (1987) Memo 994, MIT AI Laboratory, Cambridge, MA.

Horn, B.K.P. (1990) "Height and Gradient from Shading," to appear in *International Journal of Computer Vision*. Also, (1989) Memo 1105, MIT AI Laboratory, Cambridge, MA.

Horn, B.K.P. & M.J. Brooks (1986) "The Variational Approach to Shape from Shading," *Computer Vision, Graphics and Image Processing*, Vol. 33, No. 2, pp. 174–208, February. Also *Shape from Shading*, Horn, B.K.P. & M.J. Brooks (eds.), MIT Press, Cambridge, MA..

Horn, B.K.P. & M.J. Brooks (1989) *Shape from Shading*, MIT Press, Cambridge, MA.

Horn, B.K.P. & B.G. Schunck (1981) "Determining Optical Flow," *Artificial Intelligence*, Vol. 17, No. 1–3, August 1981, pp. 185–203. Also (1980) Memo 572, MIT AI Laboratory, Cambridge, MA. Also AD-A093925 available from NTIS.

Horn, B.K.P., R.J. Woodham & W.M. Silver (1978) "Determining Shape and Reflectance using Multiple Images," Memo 490, MIT AI Laboratory, Cambridge, MA. Also AD-A062380 available from NTIS.

Howard, A.D., K.R. Blasius, & J.A. Cutt (1982) "Photoclinometric Determination of the Topography of the Martian North Polar Cap," *Icarus*, Vol. 50, pp. 245–258.

Ikeuchi, K. & B.K.P. Horn (1981) "Numerical Shape from Shading and Occluding Boundaries," *Artificial Intelligence*, Vol. 17, No. 1–3, pp. 141–184, August. Also in *Computer Vision*, Brady, J.M. (ed.), North Holland. Also in *Shape from Shading*, Horn, B.K.P. & M.J. Brooks (eds.), MIT Press, Cambridge, MA..

Kirk, R.L. (1984) "A Fast Finite-Element Algorithm for Two-Dimensional Photoclinometry," *Bulletin of the American Astronomical Society*, Vol. 16, No. 3, pg. 709.

Kirk, R.L. (1987) "A Fast Finite-Element Algorithm for Two-Dimensional Photoclinometry," Part III of Ph.D. Thesis, Division of Geological and Planetary Sciences, California Institute of Technology, Pasedena, CA.

Lambiotte, J.J. & G.R. Taylor (1967) "A Photometric Technique for Deriving Slopes from Lunar Orbiter Photography," *Proceedings of the Conference on the Use of Space Systems for Planetary Geology and Geophysics*, Boston, MA, May 25–27.

Lucchitta, B.K. & N.A. Gambell (1970) "Evaluation of Photoclinometric Profile Determination," in *Analysis of Apollo 8 Photographs and Visual Observations*, NASA SP-201, National Aeronautics and Space Administration, pp. 51–59.

Malin, M.C. & G.E. Danielson (1984) "Topography on Ganymede Derived from Photoclinometry," in *Reports on Planetary Geology Program* (1983), NASA Technical Memorandum 86246, pp. 29–30.

McEwen, A.S. (1985) "Topography and Albedo of Ius Chasma, Mars," *Lunar and Planetary Science Conference* XVI, 11–15 March, pp. 528–529.

Passey, Q.R. & E.M. Shoemaker (1982) "Craters and Basins on Ganymede and Callisto: Morphological Indicators of Crustal Evolution," in Morrison, D. (ed.) *Satellites of Jupiter*, University of Arizona Press, Tucson, pp. 379–434.

Pentland, A.P. (1988) "Shape Information from Shading: A Theory about Human Perception," Technical Report 103, Vision Sciences, MIT Media Laboratory, MIT, Cambridge, MA, May.

Rindfleisch, T. (1966) "Photometric Method for Lunar Topography," *Photogrammetric Engineering*, Vol. 32, No. 2, pp. 262–277, March.

Rowan, L.C., J.F. McCauley & E.A. Holm (1971) "Lunar Terrain Mapping and Relative Roughness Analysis," U.S. Geological Survey Professional Paper 599-G, pp. 1-32.

Tyler, G.L., R.A. Simpson & H.J. Moore (1971) "Lunar Slope Distributions: Comparison of Bi-Static Radar and Photographic Results," *Journal of Geophysical Research*, Vol. 76, No. 11, pp. 2790–2795.

Szeliski, R. (1990) "Fast Shape from Shading," Internal Report, DEC Cambridge Research Center, Cambridge, MA.

Watson, K. (1968) "Photoclinometry from Spacecraft Images," U.S. Geological Survey Professional Paper 599-B, pp. 1–10.

Wildey, R.L. (1986) "Radarclinometry for the Venus Radar Mapper," *Photogrammetric Engineering and Remote Sensing*, Vol. 52, No. 1, pp. 41–50, January. Also in *Shape from Shading*, Horn, B.K.P. & M.J. Brooks (eds.), MIT Press, Cambridge, MA..

Wilhelms, D.E. (1964) "A Photometric Technique for Measurement of Lunar Slopes," in *Astrogeological Studies Annual Progress Report, Part D: Studies for Space Flight Program*. U.S. Geological Survey Open-File Report, pp. 1–12, May, NASA Catalog Number N66 35597.

Wilson, L., M.A. Brown, E.M. Parmentier & J.W. Head (1984) "Theoretical Aspects of Photoclinometry Terrain Profiling on the Galilean Satellites," in *Reports on Planetary Geology Program* (1983), NASA Technical Memorandum 86246, pp. 27–28.

Zheng Q. & R. Chellapa (1990) "A Robust Algorithm for Inferring Shape from a Single Image," Working Paper, USC, SIPI, Los Angeles, CA.